# THE TECHNOLOGY OF DISASTER RECOVERY

Sheri Atwood

October 8, 2003

VERITAS ARCHITECT NETWORK

# TABLE OF CONTENTS

## Imagining the Worst

In growing numbers, organizations are facing a radical shift in value that emphasizes knowledge over the delivery of products and services. This change, which is a result of increased commoditization, highlights the critical nature of access to an organization’s data assets.

Information technology (IT) systems that broker the delivery of business information provide enormous value, but they also introduce a vulnerability. If access to vital data is interrupted, the organization suffers.

At a minimum, outages to IT systems can cost millions of dollars in lost revenue, lost productivity, and legal issues. At the extreme, a sustained outage can threaten the viability of an organization. According to the Gartner Group, “two out of every five enterprises that experience a disaster go out of business within five years.” (Gartner, *Disaster Recovery Plans and Systems are Essential*, Robert Witty, Donna Scott, September 2001)

It no longer requires a tremendous stretch of the imagination to envision scenarios that can cripple an organization’s technology assets. Whether the result of a terrorist attack or a power failure, organizations on both the East and West coasts of the United States have had recent reminders of the critical importance of planning for disasters. No matter how unlikely it may seem today, every organization must face the near certainty of a business-wide failure of IT systems occurring at a future date. Anticipating these events and planning corrective courses of action is now a prerequisite to business success.

Preparing an organization for the unexpected is the domain of business continuity planning (BCP). The specific preparations taken by an IT group to ensure continuous access to information resources is a subset of BCP known as disaster recovery planning (DRP). This article provides an overview of the technology of disaster recovery (DR) in order to give application developers, system administrators, and line-of-business users a common reference point when discussing how best to respond to the threat of a business system outage.

## The Language of Disaster Recovery

The terminology of DRP can seem confusing and cryptic to those not involved in disaster planning on a daily basis. Many DR terms appear bewilderingly similar – high availability and disaster tolerance, for example – and others are loaded with mysterious technology-related jargon.

For application analysts and line-of-business personnel involved in the DR process, understanding the terminology used to describe DRP is essential. Proposals for safeguarding critical business systems will often be peppered with arcane and unusual phrases, and assessing the merits of a particular DR approach demands a common understanding of the language.

### Business Continuity Planning and Disaster Recovery Planning

Business continuity planning (BCP) is a procedure put in place by an organization to ensure that essential business processes continue following a disaster. The BCP takes into account the need for alternate facilities (offices, warehouses, and retail outlets) if normal business locations become inaccessible. A host of other items are included in the plan, including departmental guidelines detailing how to maintain business operations under extraordinary circumstances.

Disaster Recovery Planning (DRP) is a subset of BCP and focuses solely on the recovery of IT systems. The DR plan, the output of the DRP process, documents procedures for IT staff to follow when reestablishing business system functionality after an outage. Each business application must be cataloged, its recovery needs assessed and documented, and the importance of the application to the organization quantified to enable IT staff to prioritize the recovery process.

## Business Impact Analysis

A business impact analysis (BIA) quantifies the impact of an outage to each business system. The BIA determines what affect the loss of a specific IT system will have on an organization. For example, a failure disrupting the accounts payable system may have wide ranging consequences for cash flow, customer retention, and an organization's credit rating.

The BIA includes a risk analysis to determine the likelihood of a disruption to business applications. The probability of an event is weighed against the amount of disruption the event might cause. Findings from the BIA help the IT department determine strategies to offset the risk of an event's occurrence.

## Hotsites, Warmsites, and Coldsites

Although the specific characteristics of a hotsite, warmsite, or coldsite vary between organizations, these terms are universally used to grade the state of readiness of a remote data center.

A DR hotsite provides a fully operational computing environment that includes servers, storage, and networking equipment. Applications and data at the hotsite are closely synchronized with the primary data center. In a disaster, operational support of IT systems can be quickly switched from the primary site to the hotsite. The prompt failover of applications to a hotsite minimizes the impact of an outage on the business.

A warmsite generally refers to a data center facility with all the necessary hardware and communications equipment needed to run a business; however, the systems are not kept in a constant state of operational readiness. When a disaster is called, applications and data must be recovered at the warmsite to provide support for ongoing business operations.

A coldsite facility has no hardware, but it provides power, communications access, and an environment for hosting a computing infrastructure. Following a disaster, IT staff must recreate the data center from scratch, and considerable work is needed before the facility can host business applications.

## Disaster Tolerance

Greater awareness of the need for DRP is prompting application architects to build disaster readiness into the business systems they design. Disaster tolerance is a term used to signify a system with some ability to withstand major disruption. Several technologies are used to provide disaster tolerance, including hardware redundancy, data replication, server clustering, and remote data centers.

## High Availability Systems

The ultimate disaster tolerant system is classed as a high availability (HA) system. This configuration is designed to eliminate application downtime by using redundant hardware and networking components and specialized application and operating system software. HA systems can seamlessly route around failures in the computing infrastructure without affecting end-user access to data.

The resilience of this system is often measured in terminology borrowed from the telecommunications industry. For example, a configuration that offers 99.999% availability (also known as *five nines*), will be unavailable for no more than five minutes in any given year.

## RPO and RTO

The business impact analysis (BIA) produces two key metrics that measure a business system's ability to tolerate lost data and downtime – recovery point objective (RPO) and recovery time objective (RTO). The RPO denotes the amount of data an application can lose before an organization begins to suffer. The RTO indicates how much time the IT staff can take to bring the application back online after a disaster occurs. The unit of measure for both RPO and RTO is time, with values ranging from seconds to days or weeks. The closer an application's RPO and RTO values are to zero, the greater the organization's dependence on that particular process, and consequently the higher the priority when recovering the systems after a disaster.

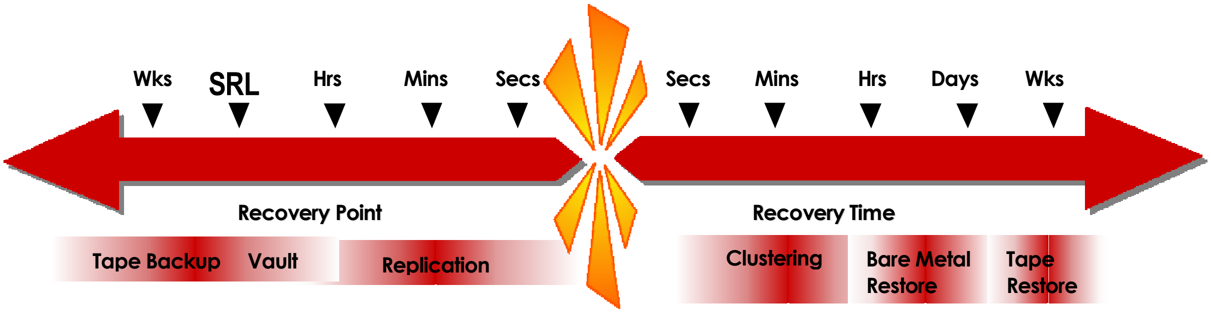

Figure 1. Translating RPO and RTO to Technologies

## Backup and Recovery

Backup refers to the process of copying information from disk to a secure storage medium (usually tape), and backup also refers to the storage medium itself. Tape backup provides one of the fundamental building blocks of a DR plan.

Determining when to run a backup is a function of the interconnectivity and interdependency of many different business applications. Backup administrators carefully coordinate the copy process to ensure the integrity of the information on tape. Recovering data from a backup involves restoring the contents of the tape to disk and then performing reconciliation processing to rectify any errant information.

## Tape Vaulting

Vaulting backup tapes guarantees that backup data does not suffer the same fate as primary data when a disaster strikes. Basic vaulting of tape data is as simple as sending backups offsite for storage at a secure location.

Electronic vaulting of tape data is an attractive option because of widely available and relatively low-cost, high-speed data networks. Rather than transport tape copies offsite manually (a notoriously error prone process), electronic vaulting systems transmit backup data over the network to tape devices located at a secure vaulting facility or alternate data center.

Automated electronic vaulting improves the prospects of successfully recovering data after a disaster. Manual errors are significantly reduced, and the end-to-end elapsed time of the backup process is shortened to allow backups to be run more frequently.

## Replication and Remote Mirroring

The terms *replication* and *remote mirroring* are frequently used interchangeably because their meanings are very similar. Replication is the process of duplicating primary data volumes over an IP connection to a storage subsystem at an alternate location. This redundant copy can then be used by business applications if access to the primary data is interrupted. The source and target of a replication are usually separated by a significant distance to safeguard data from disasters that effect a specific geographic location, such as a region-wide power outage.

Replication techniques have two principle modes: asynchronous and synchronous. With asynchronous replication, the primary and secondary data volumes are no more than a few milliseconds out of sync, so the replication is nearly real-time. With synchronous replication, the primary and secondary copies are always identical, so it provides a true real-time duplication. (See Replication.)

Like replication, remote mirroring uses redundancy to guarantee data availability. A remote mirrored data volume consists of two identical copies of the data connected by Fibre Channel. Both sides of a mirror process read and write I/Os to ensure each copy is a real-time duplicate of the other. The mirror copies are housed in separate data centers, connected over a local- or metro-area network (LAN or MAN). If a disaster interrupts access to one side of a mirror, the surviving copy will continue to service application I/O requests and therefore minimize the disruption to business users.

## Bare Metal Restore

One of the first tasks of any DR plan is to recover the backup server. The process can require a significant amount of time, because it requires the reconfiguring of server hardware at the DR site so that it is identical to the production environment. You must painstakingly apply operating systems, software patches, and site-specific customizations to create an exact copy of the production server.

Bare metal restore technologies simplify and speed backup server recovery. An image of the production backup server is generated and sent offsite with backups of business data. In a disaster, the backup server can be instantly rebuilt from the backup image, so the process of restoring business data can take place immediately.

## Clustering

Clustering uses redundancy to minimize the impact of an outage to end users. In a clustering environment, multiple application servers are configured to operate as a single computing unit. Each server is referred to as a cluster node. If an application or node fails, the impacted processing automatically fails over to a surviving node. Clustering is configured for either local failover, using mirroring technologies over a LAN or MAN, or for global outage protection, using replication techniques across a wide-area network (WAN).

# Anatomy of a Disaster

Not all unexpected outages constitute a disaster. Calling an event a disaster has a very specific meaning to IT staff, with an associated set of consequences that are carefully documented in the DRP.

## A Business System Outage

For a variety of reasons, end users will occasionally lose access to business applications. Technicians staffing an IT customer support desk are trained to triage would-be disasters. This allows quick identification of lesser problems, such as a user’s PC crashing, local network overload, or mishaps in the data center that can be quickly rectified. However, if an event is recognized as having disaster potential, a swift analysis of the available options is critical.

The DR plan dictates roles that IT staff members will assume when a disaster occurs. A team will be assembled to assess the extent of the problem and determine how the DR plan will be implemented.

Earthquakes, floods, and other high-profile catastrophic events grab media attention, but the DR plan is invoked more likely for localized incidents such as a data-center power outage, a critical hardware failure, or a software-related problem such as a virus. The RTO of each application determines how the DR team responds during the initial phases of the DRP. If an application has an RTO of one hour, and initiating the DRP for that application will take 45 minutes, the DR team has 15 minutes to determine whether a solution can be arrived at without initiating the DRP. The end of this 15-minute interval is referred to as the mandatory decision point.

## Losing a Data Center

Events that impact a data center to such an extent that it can no longer function are less ambiguous than the loss of an individual application. These events are likely to have ramifications beyond IT and will require the initiation of the BCP. The inadvertent slicing of a power line by a backhoe operator is the most innocuous of reasons for a data center outage. Earthquakes, floods, and terrorism offer complications that are much more significant. Regardless of the cause, the loss of data center facilities will result in the DR plan being invoked and the recovery of applications according to pre-arranged procedures.

## Disaster Recovery Strategies

The strategies used to protect data from loss due to a disaster must reflect the priorities of the organization. Spending $1 million to ensure the speedy recovery of a file and printer server may be overkill, but budgeting the same amount to safeguard a critical revenue generating application is likely justified. The RPO and RTO give IT administrators the information they need to identify an appropriate strategy for a particular application. These two key DRP metrics can also help verify the success of a chosen strategy during DR testing.

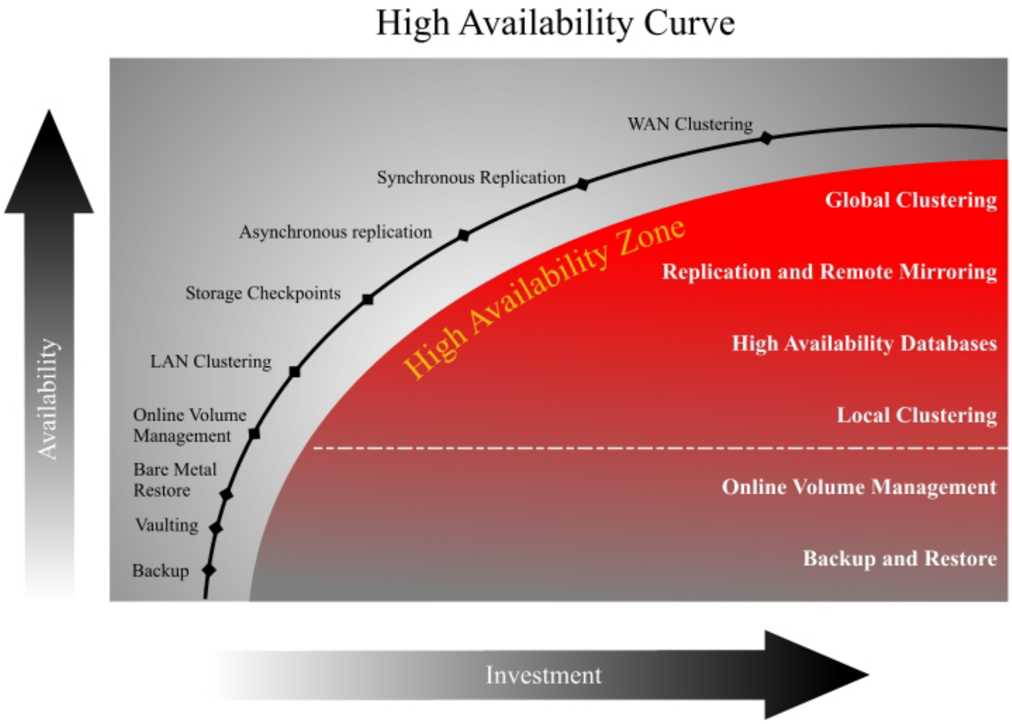

Figure 2. The High Availability Curve

### The Backup Plan

Backing up business data to tape is the traditional form of DRP employed by most IT organizations. Although the interdependence of business systems can make meeting a backup window challenging, the recovery of an application following a disaster can be relatively quick and painless once a successful backup procedure is in place.

#### Onsite Tape Storage

Tape backups provide an essential component of everyday application data protection. Organizations invest significant sums on tape-based systems, including high-speed, high-density tape devices and sophisticated, robotically operated tape libraries. Storing tapes that will be used for DR in an onsite tape library, however, cannot be considered part of a DR plan. The risk of losing DR data along with the application data is too great. Unless backup tapes are taken offsite for safe storage, the recovery of an application after a disaster is left completely to chance.

### Offsite Tape Storage

Transporting backup tapes offsite for storage at a secure location is sometimes referred to as PTAM (pick-up truck access method) tape vaulting. Whether transporting to the CFO's basement or to a disaster recovery vendor's tape vault, this method offers an appropriate means of safeguarding data for many business applications. The tapes can be retrieved after a disaster – along with catalog information detailing the name and contents of each tape – and business applications can be recovered at the DR site.

Catalog data can be restored manually at the disaster site by rebuilding the catalog. However, many organizations replicate catalog information to an offsite server so it will be online and immediately accessible in the event of a disaster. A replicated catalog will significantly decrease the amount of time it takes to restore data from tape.

The RTO of applications that rely on tape-based DR strategies is measured in days, rather than hours. The recovery process involves locating the appropriate tapes and transporting them to a DR site where they can be restored. The interdependence of application systems often complicates the recovery process, requiring a specific sequence of restores before restarting the desired application.

### Electronic Vaulting to a Hotsite

The electronic vaulting of tape backups reduces tape-based application recovery time from days to hours, making it a flexible alternative for business systems with more demanding RTOs. Using high-speed, high-bandwidth data networks, the electronic vaulting process writes backup data directly to a tape device located at a hotsite or warmsite data center. With greater scope for automation than traditional tape vaulting, electronic vaulting increases the chances of successfully restoring data. Electronic vaulting also speeds the backup process. This can allow application data to be sent offsite more frequently, perhaps avoiding the need for more expensive DR strategies.

Tape-based DRP strategies, whether using electronic vaulting or not, have an associated high risk of data loss. A backup reflects the data at a specific point in time, and any changes to the data after that point will be lost when the backup is restored. Therefore, business processes that cannot tolerate missing information require a more sophisticated approach to DRP.

## The Second Data Center

Application systems with the most demanding RPOs and RTOs need service level guarantees beyond the capabilities of tape-based DR solutions. For these strategic business applications, data replication and clustering are preferred for DRP.

Effective replication and clustering require considerable control over the DR environment. Maintaining connectivity between remote and local servers is essential for both the success of the DRP and day-to-day application performance. A successful replication or clustering strategy invariably requires investment in a second data center, significantly increasing the organization's commitment to DRP.

### Replication

Data replication reduces the window for data loss during a disaster and speeds the application recovery process. There are two basic forms of replication: asynchronous and synchronous. Asynchronous replication provides near real-time copies of production data, while synchronous replication supports full, real-time duplication of data.

With both methods, updates that are made to a primary data volume are duplicated on a secondary volume located at a remote data center. Because updates are transmitted immediately to the secondary volume, the two copies of data are kept closely synchronized. In a disaster, business applications are rerouted to the secondary volume, allowing very fast failover with little or no data loss.

These two modes of replication differ in how they manage the interface between a business application and the physical I/O it generates. Synchronous replication treats the I/O to both local and remote data volumes as a single process. Both I/Os must succeed, or fail, before the application can consider the process complete. Asynchronous replication handles the local and remote I/Os differently. Each I/O is considered an independent process, and the business application will only wait until the local write is complete before continuing.

The advantage to decoupling local and remote I/O processes is the freedom from network latency. Although latency is unlikely to be a significant factor when the primary and secondary data centers are connected over a LAN or MAN, the delay from writing to a storage volume in a data center at the other end of a WAN will almost certainly have an unacceptable impact on application performance.

The downside to I/O independence is an increased risk of data loss. During asynchronous replication processing, there will always be updates that have been successfully written to local volumes but not yet committed to disk at the remote data center. Although the difference between primary and secondary volumes may be only milliseconds, the loss of these in-flight I/Os during a disaster will be unacceptable for applications with the most rigorous RPOs.

Synchronous replication resolves the lost in-flight I/O problem, ensuring that local and remote volumes are kept perfectly matched. The impact of network latency on performance, however, will limit the application of synchronous replication to storage volumes in data centers within the same geographic vicinity.

## LAN, MAN, and WAN Clustering

The reality of modern DR methods favors a tiered approach to DRP. Clustering configurations offer IT administrators the flexibility to satisfy the most stringent RPO and RTO demands, and yet deliver a response appropriate to the type of disaster (see Figure 3).

LAN- and MAN-based clustering configurations provide significant protection from local disaster events. Deploying a LAN connected multi-node cluster in the primary data center allows failover from disaster events contained in the local environment. For example, if a data center maintains redundant power supplies and one supply fails, a clustered server configuration will fail over application I/O requests to nodes attached to the functioning power supply.

Stretching a cluster configuration between two data centers connected by a MAN provides applications with another layer of protection. If a failure impacts the primary data center, cluster management software can automatically fail application access over to data at the second data center.

LAN and MAN clusters use mirroring to synchronize data volumes between nodes. Mirroring provides synchronous, real-time data duplication, eliminating data loss and minimizing downtime from a disaster. However, mirroring has distance limitations. Beyond the range of a MAN, excessive network latency begins to have an unacceptable impact on application performance.

The geographic coverage offered by a MAN may also fail to provide the isolation necessary to protect data from disasters that affect an entire region. An organization with a primary data center in New York and a second MAN-attached data center in New Jersey, for example, would have suffered outages at both facilities during the recent multi-state power outage. With data centers separated by significant distances, organizations can avoid this type of failure. WAN clustering offers the levels of disaster protection needed to isolate data centers from geography-centric catastrophes.

WAN clustering provides the same level of automated failover as LAN and MAN clustering. However, it uses replication technology, rather than mirroring, to distribute data between nodes. Asynchronous replication removes network latency from the DR equation and allows primary and secondary copies to maintain near real-time synchronization.

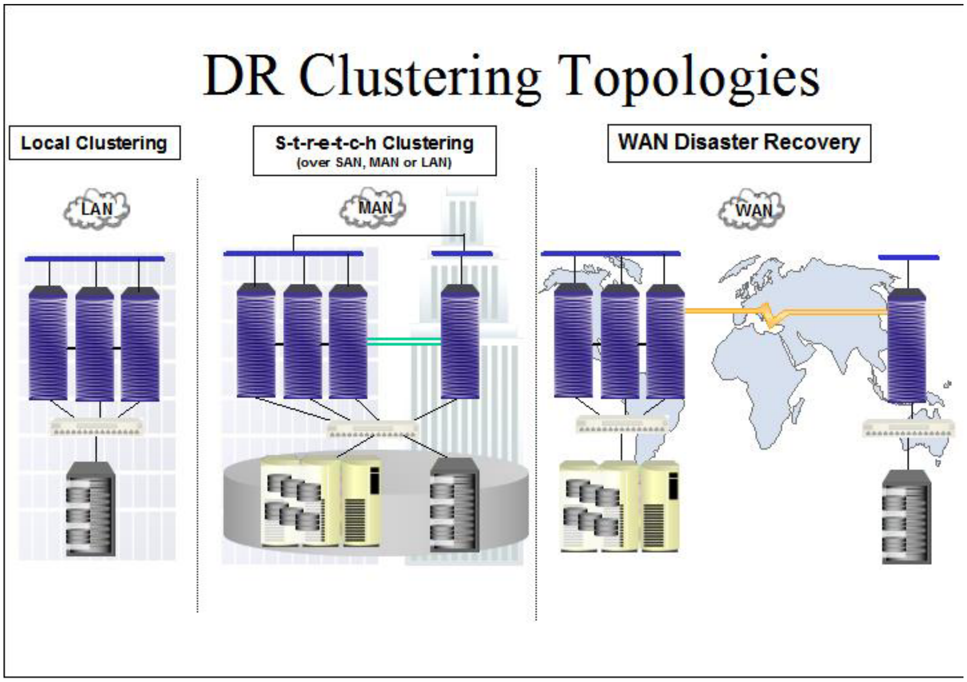

Figure 3. DR Clustering Topologies

## Summary

High-speed networks now make it possible to host copies of production data at distant locations without requiring the use of unpredictable tape-based recovery. The widespread availability, and relatively low cost, of high-bandwidth networks has allowed data replication to supplant traditional magnetic tape as the key to disaster recovery effectiveness.

The reality of modern production IT environments will dictate the use of a range of different disaster protection technologies. However, understanding the needs of each application is the first step to an effective DR plan. No matter what technologies are available, the RPO and RTO offer the most reliable guide to the level of protection that a business demands for each application.

**VERITAS Software Corporation**

Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 866-837-4827

For additional information about VERITAS Software, its products, VERITAS Architect Network, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com.